



NERSC Biology & Computing



David Skinner
Outreach, Software, Programming

Diffuse X-rays and Protein Dynamics

NERSC History



Cray 1 - 1978

Cray 2 – 1985



Cray T3E Mcurie - 1996





IBM Power3 Seaborg - 2001

1974	Founded at Livermore to support fusion research with a CDC system
1978	Cray 1 installed
1983	Expanded to support today's DOE Office of Science
1986	ESnet established at NERSC
1994	Cray T3D MPP testbed
1994 - 2000	Transitioned users from vector processing to MPP
1996	Moved to Berkeley Lab
1996	PDSF data intensive computing system for nuclear and high energy physics
1999	HPSS becomes mass storage platform
2010	Collaboration with Joint Genome Institute (JGI)

We directly support DOE's science mission

- We are the primary computing facility for DOE Office of Science
- DOE SC allocates the vast majority of the computing and storage resources at NERSC
 - Six program offices allocate their base allocations and they submit proposals for overtargets
 - Deputy Director of Science prioritizes overtarget requests
- Usage shifts as DOE priorities change



We focus on the scientific impact of our users

- 1500 journal publications per year
- 10 journal cover stories per year on average
- Simulations at NERSC were key to **2 Nobel Prizes** (2007 and 2011)
- Supernova 2011fe was caught within hours of its explosion in 2011 and telescopes from around the world were redirected to it the same night
- Data resources and services at NERSC played important roles in **two of Science Magazine's Top Ten Breakthroughs of 2012** — the discovery of the Higgs boson and the measurement of the Θ_{13} neutrino weak mixing angle
- MIT researchers developed a new approach for desalinating sea water using sheets of graphene, a one-atom-thick form of the element carbon. **Smithsonian Magazine's fifth "Surprising Scientific Milestone of 2012."**
- **Four of Science Magazine's insights of the last decade** (3 in genomics, 1 related to cosmic microwave background)





U.S. DEPARTMENT OF
ENERGY

Office of
Science



Our operational priority is providing highly available HPC resources backed by exceptional user support

- **We maintain a very high availability of resources (>90%)**
 - One large HPC system is available at all times to run large-scale simulations and solve high throughput problems
- **Our goal is to maximize the productivity of our users**
 - One-on-one consulting
 - Training (e.g., webinars)
 - Extensive use of web technologies
 - We solve or have a path to solve 80% of user tickets within 3 business days



NERSC Strategic Objectives 2013

- **Meet the ever growing computing and data needs of our users by**
 - providing usable exascale computing and storage systems
 - transitioning SC codes to execute effectively on many core architectures
 - influencing the computer industry to ensure that future systems meet the mission needs of SC
- **Increase the productivity, usability, and impact of DOE's user facilities by providing comprehensive data systems and services to store, analyze, manage, and share data from those facilities**

Mission & Extreme Data Analysis




- **NERSC Mission: Accelerate scientific discovery through high performance computing and extreme data analysis**
- **2013 Goal: Become a national exchange and analysis engine for massive and diverse scientific data sets.**
 - 4V's → Volume, Velocity, Variety, Veracity
 - High performance data sharing capabilities with an integrated vision for how scientists will consume, analyze annotate, crowd-source, and publish large data sets.
 - New algorithms and software for traversing, searching, indexing, distributing and thinking about BigData
- **Photon Science**
 - Emerging needs for transforming data methods for DOE instruments
 - Need for a data facility
- **Overlap : Making BigData deliver for science**

NERSC's Data Directions


- **Become a data science hub as well as a computing center**
- **Field cutting edge HPC storage and data rechnologies**
- **Bring BigData solutions to science partners who are ready**
 - Bioinformatics, bioimaging and health (JGI, NGBI LDRD, UCSF/NIH)
 - Photon science, Synchrotrons (ALS, SLAC, NGLS)
 - Materials Genomics, Manufacturing (MatProj, BES, JCESR, CA)
- **Science gateways**
 - Bring advanced computing and data technology to world's scientists
 - Database-driven workflows and storage
 - Scalable structured and unstructured object stores
 - Software solutions to traverse massive data for search or analysis
 - Sophisticated web-based gateways to interact with and leverage data
 - Comprehensive scientific data curation beyond simple archiving

Data is Reshaping HPC

- The observational dataset for the Large Synoptic Survey Telescope will be ~100 PB
- The Daya Bay project will require simulations which will use over 128 PB of aggregate memory
- By 2017 ATLAS/CMS will have generated 190 PB
- Light Source Data Projections:
 - 2009: 65 TB/yr
 - 2011: 312 TB/yr
 - 2013: 1.9 PB /yr
 - EB in 2021?
 - NGLS is expected to generate data at a Terabit per second




Big Data Analysis @ **NERSC**



Light Sources

- Many detectors on > Moore's Law curve
- Data volumes rendering previous operational models obsolete




High Energy Physics

- LHC Experiments produce and distributed petabytes of data/year
- Peak data rates increase 3-5x over 5 years, TB of data per night



Genomics

- Sequencer data volume increasing 12x over next 3 years
- Sequencer cost decreasing by 10 over same period



Computing

- Simulations at Scale and at High Volume already produce Petabytes of data and datasets will grow to Exabytes by the end of the decade
- Significant challenges in data management, analysis and networks



U.S. DEPARTMENT OF
ENERGY

Office of
Science




PS-II experiment, ESnet-scale data flows

All NERSC
Traffic



Photosystem II
X-Ray Study



U.S. DEPARTMENT OF
ENERGY

Office of
Science




LCLS Diffract & Destroy

Photo-system II, Nick Sauter (LBL): A single experiment in 2013 generated **120TB** over four days. Data was relayed to NERSC via Esnet, analysis required **135K CPU hours** of computing.





LCLS-II will require **> 100M CPU hours** per experiment. Higher resolution and advanced image analysis could grow computational complexity. Some algorithms scale an $M \cdot N \log N$ for M images of N pixels.



Traffic split by : 'Autonomous System (origin)'

nersc-SLAC:3671




U.S. DEPARTMENT OF
ENERGY

Office of
Science




Strategy for Big Data Era Lightsources (I)

- **Accommodate Big Data from high-speed detectors**
 - Advanced cameras and DAQs, make detectors smarter
 - Scalable data formats (HDF) and parallel IO (GPFS,Lustre)
- **Automated Big Data pipelines**
 - Move data quickly (ESnet) to reliable persistent storage
 - Pub/sub model for both prompt and off-line needs
 - Centered at the world's fastest science network
- **Scalable data-centric computing partnerships**
 - Lessons from the pilots: detector → network → HPC → community
 - Organized, scalable, on-demand data analysis capabilities
- **Extend science impact via the web**
 - Working with Big Data remotely. Think google, not FedEx.
 - Science gateways broaden access to Lightsource & HPC capabilities.
 - Massively collaborative shared/re-usable analysis grows science impact from data.



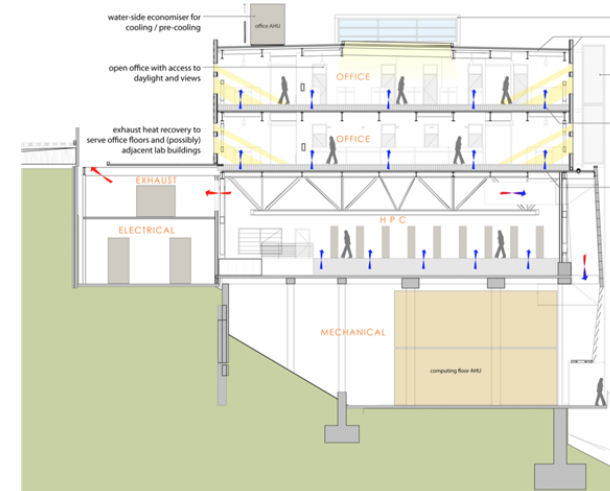
Strategy for Big Data Era Lightsources (II)

- **Meet the enemy** —————→
 - The pixel is not a great data model
- **Cameras need to get smarter**
 - LHC w/o triggers = 1PB/sec
 - LHC w/ triggers = 1PB/year
 - Pushing analysis upstream will require new R&D.
 - Don't over leverage LHC strategy
 - Infrastructure for smaller teams, w/ shorter deadlines. Think Agile.
 - Avoid the grid middleware trap, fast forward to the web.
 - Avoid political computing solutions. There's a reason google doesn't have a data center in every town. Scale works.
- **New data analysis capabilities**
 - Analysis means bandwidth & computing, high-speed data mashups
 - NERSC has > 30 years of DOE science data, production focus
 - Powerful web interfaces can share this capability with the world



We are deploying the CRT facility to meet the ever growing computing and data needs of our users

- **Four story, 140,000 GSF**
 - Two 20Ksf office floors, 300 offices
 - 20K -> 29Ksf HPC floor
 - Mechanical floor
- **42MW to building**
 - 12.5MW initially provisioned
 - WAPA power: Green hydro
- **Energy efficient**
 - Year-round free air and water cooling
 - PUE < 1.1
 - LEED Gold
- **Occupancy Early 2015**



Addressing the science “Big Data” challenge requires new math, algorithms, software, computing systems

- **Research: Perform basic research in data science**
 - **Math and Algorithms:** Statistics & Machine Learning, Image analysis, Graph analytics
 - **Computing Systems:** Develop/evaluate new hardware, programming, and software techniques;
- **Software**
 - **Partnerships:** Leverage Lab staff and culture to develop usable robust tools for data movement, provenance, analysis; create end-to-end scientific workflows
 - **Cross-lab organization:** to encourage re-use and sharing
- **Facilities: Drive and deploy the best technology for data science in collaboration with industry**
 - **Networking:** for bandwidth, flexibility, and to federate facilities
 - **Systems:** for data storage, collaborative analysis, dissemination

Looking to the future....

- **Many computing challenges remain the same**
 - power, concurrency, data movement, programmability
- **Progress depends on close collaboration between computing, networking and computer science**
- **Large memory, processing near memory, active storage, scalable data access, machine learning**
- **Human in the loop requires fast turn-around**
- **Improving cycle time and performance of device yields tangibly measured science and cost benefits**

NERSC and Esnet: WAN data trends

**Daily WAN traffic in/out
of NERSC over the last
decade**

**Roughly 10x ↑
from 2011 to 2016**

Automated data pipelines ,
Esnet

Large scale image processing ,
NERSC

Community access to data and
analysis , gateways

NERSC Daily WAN Traffic since 2001

